

A FAULT-TREE-BASED IMPORTANCE SAMPLING STRATEGY FOR PIECEWISE DETERMINISTIC MARKOV PROCESSES

Guillaume Chennetier^{1,2*}, Hassane Chraïbi², Anne Dutfoy² and Josselin Garnier¹

¹CMAP, CNRS, Ecole polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France.
e-mail: guillaume.chennetier@polytechnique.edu, josselin.garnier@polytechnique.edu

²EDF Lab Paris-Saclay, Boulevard Gaspard Monge, 91120 Palaiseau, France.
e-mail: guillaume.chennetier@edf.fr, anne.dutfoy@edf.fr, hassane.chraïbi@edf.fr

Abstract. *Piecewise deterministic Markov processes (PDMPs) can be used to model complex dynamical industrial systems. The counterpart of this modeling capability is their simulation cost, which makes reliability assessment untractable with standard Monte Carlo methods. Indeed, the failure of a complex system is a rare event and estimating the probability of its occurrence using a Monte Carlo method requires the simulation of a very large number of trajectories of the underlying process. A significant variance reduction can be obtained with a well-calibrated importance sampling method. It is known that the optimal distribution for importance sampling depends explicitly on the committor function of the PDMP. Fault tree analysis offers us elegant tools to approximate this committor function. We present an adaptive importance sampling (AIS) method based on a cross-entropy (CE) procedure for sequentially refining the approximation of the committor function. The method is tested on a system from the nuclear industry: the spent fuel pool.*

Keywords: Rare event simulation, Importance sampling, Piecewise deterministic Markov process, Reliability assessment, Fault tree analysis.

1 INTRODUCTION

Our work aims to evaluate the reliability of complex dynamic industrial systems that play a crucial role in nuclear power plants and dams. These systems can experience critical failure when one or more continuous physical variables, such as temperature or pressure, exceed a critical threshold. This threshold can only be reached after the deterioration of some key sets of system components, which have discrete states that change due to random point events such as failures, repairs, and control mechanisms. The continuous physical variables evolve according to deterministic differential equations that depend on the state of the system's components. This hybrid behavior leads us to model the system's state over time using a piecewise deterministic Markov process (PDMP) [1].

However, assessing the reliability of such systems is challenging. Experimental data may not be available, and the complexity of a PDMP makes it impossible to compute the probability of system failure. Numerical simulations are therefore often used. But in industries, critical failures are rare events, and the estimation of a rare event probability requires a large number of simulations when using crude Monte Carlo methods (CMC). Furthermore, solving the differential equations involved in PDMP simulations can be computationally intensive, which further increases simulation time. To reduce the number of simulations, we propose an importance sampling method. The role played by the committor function of the process in the optimal importance sampling distribution has been well described in [3]. The committor function returns the probability that a PDMP trajectory realizes the rare event of interest (in our case, system failure) given its state at a given instant. The key point of our method is to propose a relevant approximation of this committor function to build a parametric family of importance distributions.

A good approximation of the committor function should measure in some sense the level of degradation of the system. We use a reliability methodology called fault tree analysis [5] to determine the minimal path sets (MPS) of the system, which are the sets of components that ensure the proper functioning of the system. Each MPS can be seen as a barrier that the process must cross to reach system failure. We propose a parametric family of approximations of the committor function based on the MPS decomposition of the system, along with the corresponding family of importance distributions. We use the cross-entropy procedure [2] to sequentially determine the best representative of this family of importance distributions. By recycling past samples, we estimate the probability of system failure without additional draws while preserving the consistency and asymptotic normality of the final estimator [6]. We obtain, in practice, a variance reduction of the order of 10,000 compared to a crude Monte Carlo method on a test case from the nuclear industry: the spent fuel pool [7].

The paper is organized as follows. Section 2 provides the main definitions and results on PDMPs and presents a typical industrial system, the spent fuel pool, as an example. Section 3 explains the role of the committor function in the optimal distribution for importance sampling. Section 4 describes the reliability notion of minimal path sets from fault tree analysis and how we use it to build a family of approximations of the committor function. Our cross-entropy algorithm with recycling of past samples is presented in section 5, along with the theorem ensuring convergence and asymptotic normality of the estimator produced. In section 6, we test the method on the spent fuel pool and compare it for different sample sizes to a CMC method.

We conclude in section 7.

2 PIECEWISE DETERMINISTIC MARKOV PROCESS

The state of the PDMP at time $t \geq 0$ is denoted by $Z_t = (X_t, M_t) \in E$ where $X_t \in \mathbb{X} \subset \mathbb{R}^{d_x}$ is the *position* and $M_t \in \mathbb{M}$ is the *mode* of the PDMP (with \mathbb{M} finite or countable and $E = \mathbb{X} \times \mathbb{M}$). Let $t_{\max} > 0$, we note from now $\mathcal{Z} := (Z_t)_{t \in [0, t_{\max}]}$ a PDMP trajectory of duration t_{\max} . The behavior of the process is fully characterized by three functions:

1. The *flow* Φ that gives the deterministic trajectory.
2. The *jump intensity* λ that gives the distribution of the times of the random jumps.
3. The transition *jump kernel* \mathcal{K} that gives the distribution of the location of the process after a jump.

Flow function Φ . The position X_t of the PDMP follows a partial differential equation parametrized by the mode $m \in \mathbb{M}$ which remains constant between two jumps. The solution function of this partial differential equation is noted ϕ_m . Starting from a state $z = (x, m) \in E$, the flow function $\Phi_z = \Phi_{x,m} : t \mapsto (\phi_m(x, t), m)$ indicates thus in the absence of jump the state of the process after a time $t > 0$. Solving the PDEs is often the most expensive step in simulating the PDMPs that interest us, as it may necessitate the use of powerful physical computer codes.

If the state space has boundaries (noted ∂E), the process jumps when it reaches them. For any starting state $z \in E$ we define $t_z^\partial := \inf\{t > 0 : \Phi_z(t) \in \partial E\} \in [0, +\infty]$ the deterministic hitting time of ∂E .

Jump intensity λ . The process can also jump at random times whose distribution is characterized by a jump intensity $\lambda : E \rightarrow \mathbb{R}_+$. The more weight λ gives to a state $z \in E$, the more likely the process is to jump when it passes through this state. Assuming the process is in state $z \in E$ at time $s > 0$, we note T_z the random waiting time before the next jump starting from z . Analogous to the jump times of an inhomogeneous Poisson process, its survival function is given by:

$$\mathbb{P}(T_z > t \mid Z_s = z) = \mathbb{1}_{t < t_z^\partial} \exp\left(-\int_0^t \lambda(\Phi_z(u)) \, du\right). \quad (1)$$

Jump kernel \mathcal{K} . When the process jumps from a state $z^- \in E$, the arrival state is chosen randomly among the other states according to the jump kernel \mathcal{K}_{z^-} of probability density function $z \mapsto K(z^-, z)$ according to a reference measure ν_{z^-} on E . By convenience, the jump kernel will also refer to its density K .

Probability density function of a PDMP trajectory. Let \mathcal{E} denote the set of PDMP trajectories with values in E that are feasible according to (Φ, λ, K) . The reference measure ζ on \mathcal{E} to construct the probability density function of a PDMP trajectory has been made explicit in [3].

Let $\mathcal{Z} = (Z_t)_{t \in [0, t_{\max}]} \in \mathcal{E}$ be a PDMP trajectory and $\mathbf{n}(\mathcal{Z})$ its jump count (number of events occurring before time t_{\max}). Let z_0 be the initial state of the trajectory, t_0 the waiting time before the first jump, and for $k = 1, \dots, \mathbf{n}(\mathcal{Z}) - 1$ we note z_k the state of the process after the k -th

jump, t_k the waiting time between the k -th and the $(k + 1)$ -th jump and finally $t_{\mathbf{n}(\mathcal{Z})} = t_{z_{\mathbf{n}(\mathcal{Z})}}^\partial = t_{\max} - \sum_{k=0}^{\mathbf{n}(\mathcal{Z})-1} t_k$ the waiting time between the last jump and the end of simulation at time t_{\max} . The probability density function π of the PDMP evaluated at the trajectory \mathcal{Z} is then:

$$\pi(\mathcal{Z}) = \prod_{k=0}^{\mathbf{n}(\mathcal{Z})} [\lambda(\Phi_{z_k}(t_k))] \mathbb{1}_{t_k < t_{z_k}^\partial} \times \exp \left[- \int_0^{t_k} \lambda(\Phi_{z_k}(u)) du \right] \times \prod_{k=0}^{\mathbf{n}(\mathcal{Z})-1} K(\Phi_{z_k}(t_k), z_{k+1}). \quad (2)$$

It is worth noting that evaluating the density at a previously simulated trajectory is computationally inexpensive since there is no need to recompute the flow. From here on, we assume that E and Φ are fixed. The probability density function π of the PDMP, and consequently its probability distribution, can be described by the couple (λ, K) which captures all the stochastic properties of the process.

Modeling industrial systems: the spent fuel pool case. We use the spent nuclear fuel pool [7] as an example to demonstrate how the formal definition of a PDMP can be applied to an industrial system. The pool contains fuel at the bottom, and the water in the pool is heated by the fuel. If a critical volume of water evaporates, the system fails. To prevent this, the system uses sealed circuits connected by heat exchangers to transfer heat from an external cold water source to the pool water. The system is also connected to a general power supply. In case of any issues, there are two other identical lines, a backup electrical generator for each of the three lines, and a second water source that only the third line can access. There are fifteen components in total, including the four electrical generators, the three circuits for each of the three lines, and the two water source accesses. Each component has its own failure rate and repair rate, as given in table 2. The system's operation is depicted in fig. 1.




Figure 1: Spent fuel pool representation.

The position of the PDMP is the temperature and the level of the water in the pool. The mode is the status of each component (active: 1, inactive: 0 or broken: -1). The state space

| Physical parameters | Value | Description |
|---------------------|--|---|
| r | $2.106 \times 10^{10} \text{ J} \cdot \text{h}^{-1}$ | Residual power of the fuel. |
| C | $4180 \text{ J} \cdot \text{kg}^{-1} \cdot \text{°K}^{-1}$ | Mass heat capacity. |
| ρ | $990 \text{ kg} \cdot \text{m}^{-3}$ | Density of the water. |
| A | 77 m^2 | Area of the pool. |
| $x_S^{(1)}$ | 15 °C | Temperature of the water sources. |
| Q | $550 \text{ m}^3 \text{ h}^{-1}$ | The debit water. |
| l | $2.257 \times 10^6 \text{ J} \cdot \text{kg}^{-1}$ | Latent heat of vaporization. |
| t_{\max} | 3600 h | Duration of the mission. |
| $x_0^{(2)}$ | 19 m | Initial level of water in the pool. |
| $x_{\min}^{(2)}$ | 16 m | Critical threshold of the level of water in the pool. |

 Table 1: **Physical parameters** of the SFP. Values taken from [7].

is therefore $\mathbb{R}^2 \times \{-1, 0, 1\}^{15}$. We note $X_t = (X_t^{(1)}, X_t^{(2)})$ the position of the process at time $t \geq 0$ with $X_t^{(1)}$ the temperature of the water in the pool in °C and $X_t^{(2)}$ the water level in the pool in meters (m). The evolution of these variables is described by the differential equations:

$$\frac{dX_t^{(1)}}{dt} = \mathbb{1}_{X_t^{(1)} < 100} \times \frac{r + \rho C Q (X_t^{(1)} - x_S^{(1)}) \mathbb{1}_{M_t \notin \mathbb{M}_D}}{\rho C A X_t^{(2)}} \quad (3)$$

$$\frac{dX_t^{(2)}}{dt} = -\mathbb{1}_{X_t^{(1)} = 100} \times \frac{r}{\rho C A l} \quad (4)$$

where the physical parameters are given in the table 1 and $\mathbb{M}_{\mathcal{D}}$ is the subset of every mode in \mathbb{M} such that the water in the pool is no longer cooled (for example any mode such that the four generators are broken belongs to $\mathbb{M}_{\mathcal{D}}$). As soon as the water is no longer cooled, there is a first deterministic hitting time for the water to reach 100 °C and then a second for the water level in the pool to reach the critical threshold.

Under distribution π_0 , each component c_j has a jump rate $\lambda_0^{(j)}$ which depends on its status and on the values of the physical variables of the system. The jump intensity of the PDMP in a state $z \in E$ is the sum of the jump rates of the components in state z : $\lambda_0(z) = \sum_{j=1}^{15} \lambda_0^{(j)}(z)$. At each jump from state z^- , a component c_j is randomly selected with probability $\lambda_0^{(j)}(z^-)/\lambda_0(z^-)$ and changes status (it is repaired if it was down, and fails otherwise). The system automatically reconfigures itself by enabling or disabling components so that only the necessary components are active. Here, the higher the water temperature, the more often the circuits of the three lines will break and the longer they will take to repair.

Supplementary examples of the use of PDMPs to model industrial systems can be found in [1].

3 IMPORTANCE SAMPLING FOR PDMPs

We are looking for the probability of failure of an industrial system whose operation is modeled by the trajectory of a PDMP. Let us note \mathcal{D} the subset of faulty trajectories of \mathcal{E} . Our objective is to estimate $P := \mathbb{P}_{\pi_0}(\mathcal{Z} \in \mathcal{D})$ where π_0 is the nominal distribution of the process with jump intensity λ_0 and jump kernel K_0 (in practice given by the jump rates of the system components).

| Component | Marginal jump intensity $\lambda_0^{(j)}$ for $j = 1, \dots, 15$ | | |
|-----------------------------------|--|-------------------------------------|------------------------------------|
| | when broken | when inactive | when active |
| c_i | | | |
| $c_1 = G_0$ | $4 \cdot 10^{-2}$ | $4 \cdot 10^{-6}$ | $6 \cdot 10^{-6}$ |
| $c_{i+1} = G_i, i = 1, 2, 3$ | $8 \cdot 10^{-2}$ | $2 \cdot 10^{-6}$ | $30 \cdot 10^{-6}$ |
| $c_5 = S_1$ | $1 \cdot 10^{-2}$ | $4 \cdot 10^{-6}$ | $20 \cdot 10^{-6}$ |
| $c_6 = S_2$ | $3 \cdot 10^{-2}$ | $1 \cdot 10^{-6}$ | $5 \cdot 10^{-6}$ |
| $c_{6+i} = L_{i,1}, i = 1, 2, 3$ | $(6 - 0.03X_t^{(1)}) \cdot 10^{-2}$ | $(1 + 0.05X_t^{(1)}) \cdot 10^{-6}$ | $(3 + 0.1X_t^{(1)}) \cdot 10^{-6}$ |
| $c_{9+i} = L_{i,2}, i = 1, 2, 3$ | $(6 - 0.03X_t^{(1)}) \cdot 10^{-2}$ | $(1 + 0.05X_t^{(1)}) \cdot 10^{-6}$ | $(3 + 0.1X_t^{(1)}) \cdot 10^{-6}$ |
| $c_{12+i} = L_{i,3}, i = 1, 2, 3$ | $(6 - 0.03X_t^{(1)}) \cdot 10^{-2}$ | $(1 + 0.05X_t^{(1)}) \cdot 10^{-6}$ | $(3 + 0.1X_t^{(1)}) \cdot 10^{-6}$ |

 Table 2: **Marginal jump intensity** of each component of the spent fuel pool.

A crude Monte Carlo (CMC) provides a natural, unbiased and consistent estimator of P .

$$\widehat{P}_{\text{CMC}} := \frac{1}{N} \sum_{k=1}^N \mathbb{1}_{\mathcal{Z}_k \in \mathcal{D}} \quad \text{with } \mathcal{Z}_1, \dots, \mathcal{Z}_N \stackrel{\text{i.i.d.}}{\sim} \pi_0. \quad (5)$$

Since it takes an average of N simulations to observe one realization of an event of probability $1/N$, it is clear that a CMC method is not well suited to the estimation of very small probabilities. More precisely, the smaller P is, the larger the variance of \widehat{P}_{CMC} compared to its mean. We will therefore use a variance reduction method.

Classical importance sampling. Importance sampling is one the most popular variance reduction methods for rare event simulation. The idea is to generate the random variable (or in our case, the random process) not according to the original distribution π_0 but according to an auxiliary distribution q more likely to realize the rare event and we then correct the bias by introducing the likelihood ratio $\frac{\pi_0}{q}$ in the probability estimator. We assume that $\text{supp}(q) \subset \text{supp}(\pi_0)$ (if a trajectory is feasible with q , it is feasible with π_0 since they both are densities on \mathcal{E}).

$$\mathbb{P}_{\pi_0}(\mathcal{Z} \in \mathcal{D}) = \mathbb{E}_{\pi_0}[\mathbb{1}_{\mathcal{Z} \in \mathcal{D}}] = \mathbb{E}_q \left[\mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \frac{\pi_0(\mathcal{Z})}{q(\mathcal{Z})} \right]. \quad (6)$$

The importance sampling estimator is simply the CMC estimator of the expectation in the right-hand side in eq. (6). It is therefore also a consistent and unbiased estimator of P .

$$\widehat{P}_N^{\text{IS}} = \frac{1}{N} \sum_{k=1}^N \mathbb{1}_{\mathcal{Z}_k \in \mathcal{D}} \frac{\pi_0(\mathcal{Z}_k)}{q(\mathcal{Z}_k)} \quad \text{with } \mathcal{Z}_1, \dots, \mathcal{Z}_N \stackrel{\text{i.i.d.}}{\sim} q. \quad (7)$$

The variance of this estimator relies on the choice of q . The optimal density $q_{\text{opt}} : \mathcal{Z} \mapsto \frac{1}{P} \mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \pi_0(\mathcal{Z}) \propto \pi_0(\mathcal{Z} \mid \mathcal{Z} \in \mathcal{D})$ produces an estimator with zero variance but on the one hand its normalization constant is precisely the quantity we want to estimate and on the other hand we are not supposed to know how to simulate \mathcal{Z} in \mathcal{D} proportionally to the nominal density π_0 .

Committor function. This optimal formulation can be further refined in the case of PDMPs. We have seen that choosing an importance distribution q is equivalent to choosing a couple (λ, K) . The optimal couple $(\lambda_{\text{opt}}, K_{\text{opt}})$ corresponding to q_{opt} depends on U_{opt} the *committor function* of the process. The committor function is a key concept in transition phase theory, and it is traditionally used in rare event simulation to define optimal strategies for dynamic methods,

such as multiple splitting, rather than importance sampling. In our terms, the committor function associated with a stochastic process \mathcal{Z} and an event $\{\mathcal{Z} \in \mathcal{D}\}$ gives the probability that the trajectory \mathcal{Z} will realize the event $\{\mathcal{Z} \in \mathcal{D}\}$ before a given time t_{\max} , given the current state of the trajectory. Formally, for a state $z \in E$ and a time $s > 0$, the committor function is defined as:

$$U_{\text{opt}}(z, s) := \mathbb{E}_{\pi_0} [\mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \mid Z_s = z]. \quad (8)$$

To lighten the next equations, we also define a variant of the committor function, denoted U_{opt}^- , which returns the probability of the event given the current state of the process, with the additional assumption that the PDMP is about to jump immediately. Specifically, for any departure state $z^- \in E$ and any time $s > 0$, $U_{\text{opt}}^-(z^-, s)$ is defined as:

$$U_{\text{opt}}^-(z^-, s) := \int_E U_{\text{opt}}(z, s) K(z^-, z) d\nu_{z^-}(z). \quad (9)$$

Optimal importance sampling for PDMPs. It was shown in [3] that the optimal jump intensity λ_{opt} and the optimal jump kernel K_{opt} that produce a zero variance importance sampling estimator are given by

$$\lambda_{\text{opt}}(\Phi_z(t); s) = \lambda_0(\Phi_z(t)) \times \frac{U_{\text{opt}}^-(\Phi_z(t), s+t)}{U_{\text{opt}}(\Phi_z(t), s+t)}, \quad (10)$$

$$K_{\text{opt}}(z^-, z; s) = K_0(z^-, z) \times \frac{U_{\text{opt}}(z, s)}{U_{\text{opt}}^-(z^-, s)}. \quad (11)$$

These equations have a simple interpretation.

1. If the process is k times more likely to realize the event by jumping now in state z than later, then the optimal jump intensity λ_{opt} evaluated at z should be k times larger than the nominal jump intensity λ_0 at this point.
2. And if the process is k times more likely to realize the event by jumping now in state z^- to a specific state z than by jumping randomly according to K_0 , then the optimal jump kernel K_{opt} should be k times larger than K_0 at these points.

Therefore, knowing U_{opt} is sufficient to construct q_{opt} and thus an optimal estimator of P . In practice, U_{opt} is not known, but it can be approximated. The cross-entropy method (presented in section 5) allows us to approximate U_{opt} using a parametric family of approximations, rather than restricting ourselves to a single approximation. The best candidate in the family is sequentially updated over the course of simulations. The challenge is thus to determine a good family of proxy candidates for U_{opt} .

4 FAULT-TREE-BASED IMPORTANCE SAMPLING

The first intuition that we can have on the committor function is that it indicates the level of degradation of the system. To approximate the committor function, a family of functions that increases with the number of broken components has been proposed in [3]. This approach has shown promising results on simple systems, such as a small group of components in parallel. Our goal is to extend this approach to large, complex systems.

In a complex industrial system, the impact of component failures can vary greatly, and the system's reliability depends on redundant safety measures and the specific functions of different

component groups. Therefore, for the system to fail, either an implausibly large number of components must fail (resulting in unstable likelihood ratios), or the right components must fail at the right time.

Fault tree analysis. Fault tree analysis [5] is a set of methods designed to identify the multiple possible causes of system failure. The tree allows, in the form of combinations of elementary events and logic gates, to represent the possible scenarios backwards, from the system failure to its initial state.

The current assumption is that the system is coherent, meaning that a component failure cannot prevent the system failure nor can the repair of one component cause it. Using the fault tree, we can easily determine the unique decomposition of the system into *minimal path sets* (MPS). A path set is a group of components that, if all are functioning, ensure the proper operation of the system. A path set is minimal if it contains no other path set. Conversely, breaking at least one component in each MPS ensures system failure. Each MPS can be thought of as a barrier that must be overcome in order to cause system failure. A function that increases with the number of MPS with at least one broken component can therefore provide a useful indication of how close the system is to failure.

The eight MPS of the spent fuel pool can be visualized on the series/parallel diagram in fig. 2.




Figure 2: **Decomposition of the spent fuel pool.** MPS are vertical combinations of components (e.g. G_0 , S_1 , $L_{1,1}$, $L_{2,1}$, $L_{3,1}$).

Family of approximations of U_{opt} . We are now looking for a parametric family of functions that increase with the number of MPS with at least one broken component. Let α be the parameter vector of this family and $\mathbb{A} \subset \mathbb{R}^{d_\alpha}$ be the parameter space associated with it. Inspired by

the formula proposed in [3], we suggest:

$$U_{\alpha}(z, s) := \exp \left\{ \left(\sum_{i=1}^{d_{\text{MPS}}} \alpha_i \mathbb{1}_{\beta_z \geq i} \right)^2 \right\}, \quad \alpha = (\alpha_i)_{i=1}^{d_{\text{MPS}}}, \quad (12)$$

where d_{MPS} is the number of MPS of the system and β_z is the number of MPS with at least one broken component in the state $z \in E$. The form $x \mapsto \exp(x^2)$ guarantees that the ratios U_{α}^-/U_{α} are increasing in β_z . Two comments on this family of functions:

- When d_{MPS} is large, the vector α contains many parameters to optimize. It is possible to restrict ourselves to a smaller vector $\tilde{\alpha}$ by imposing the equality of several coordinates. For example, to get a vector $\alpha \in \mathbb{R}^{d_{\alpha}}$ from $\tilde{\alpha} \in \mathbb{R}^{d_{\tilde{\alpha}}}$ with $d_{\alpha} = k \times d_{\tilde{\alpha}}$ we impose $\alpha_i = \tilde{\alpha}_{\lfloor \frac{i-1}{k} \rfloor + 1}$ for $i = 1, \dots, d_{\alpha}$.
- $U_{\alpha} \geq 1$ whereas U_{opt} is a probability. However, U_{opt} only appears in the ratios $U_{\text{opt}}^-/U_{\text{opt}}$ or its inverse, so we only seek to approximate it up to a multiplicative constant.

As for U_{opt} , we define respectively U_{α}^- , λ_{α} and K_{α} by:

$$U_{\alpha}^-(z^-, s) = \int_E U_{\alpha}(z, s) K(z^-, z) d\nu_{z^-}(z), \quad (13)$$

$$\lambda_{\alpha}(\Phi_z(t); s) = \lambda_0(\Phi_z(t)) \times \frac{U_{\alpha}^-(\Phi_z(t), s+t)}{U_{\alpha}(\Phi_z(t), s+t)}, \quad (14)$$

$$K_{\alpha}(z^-, z; s) = K_0(z^-, z) \times \frac{U_{\alpha}(z, s)}{U_{\alpha}^-(z^-, s)}. \quad (15)$$

We can therefore construct the family of importance distributions $(q_{\alpha})_{\alpha \in \mathbb{A}}$ on \mathcal{E} from the importance jump intensity and importance jump kernel $(\lambda_{\alpha}, K_{\alpha})_{\alpha \in \mathbb{A}}$.

5 ADAPTIVE CROSS-ENTROPY PROCEDURE

We are looking for a candidate within the family $(q_{\alpha})_{\alpha \in \mathbb{A}}$ as close as possible to the target distribution q_{opt} in the sense of the Kullback-Leibler divergence:

$$\begin{aligned} \arg \min_{\alpha \in \mathbb{A}} \mathcal{D}_{\text{KL}}(q_{\text{opt}} \| q_{\alpha}) &= \arg \min_{\alpha \in \mathbb{A}} \mathbb{E}_{q_{\text{opt}}} \left[\log \left(\frac{q_{\text{opt}}(\mathcal{Z})}{q_{\alpha}(\mathcal{Z})} \right) \right] \\ &= \arg \min_{\alpha \in \mathbb{A}} \mathbb{E}_{\pi_0} [-\mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \log(q_{\alpha}(\mathcal{Z}))] \end{aligned} \quad (16)$$

We used $q_{\text{opt}}(\mathcal{Z}) = \frac{1}{P} \mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \pi_0(\mathcal{Z})$ to obtain the 2nd equality. The sequential minimization of this last expectation is called cross-entropy procedure [2].

Sequential algorithm. The function $\alpha \mapsto \mathbb{E}_{\pi_0} [-\mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \log(q_{\alpha}(\mathcal{Z}))]$ is estimated by importance sampling under an initial instrumental distribution $q_{\alpha^{(1)}}$, we determine $\alpha^{(2)}$ which minimizes this estimate and we repeat the scheme. To save the simulation budget, we reuse at each iteration all the trajectories already drawn to perform the minimization step. Similarly, all the trajectories generated during the algorithm are recycled to produce the final estimator of P . To summarize, at iteration $l \in \mathbb{N}^*$:

1. **Simulation phase.** Generate n_l trajectories $\mathcal{Z}_1^{(l)}, \dots, \mathcal{Z}_{n_l}^{(l)} \stackrel{\text{i.i.d.}}{\sim} q_{\alpha^{(l)}}$.
2. **Optimization phase.** Update the parameter of the instrumental distribution with the l last samples drawn $\left(\mathcal{Z}_k^{(1)}\right)_{k=1}^{n_1}, \dots, \left(\mathcal{Z}_k^{(l)}\right)_{k=1}^{n_l}$:

$$\alpha^{(l+1)} = \arg \min_{\alpha \in \mathbb{A}} \left\{ - \sum_{r=1}^l \sum_{k=1}^{n_r} \mathbb{1}_{\mathcal{Z}_k^{(r)} \in \mathcal{D}} \frac{\pi_0(\mathcal{Z}_k^{(r)})}{q_{\alpha^{(r)}}(\mathcal{Z}_k^{(r)})} \log \left[q_{\alpha}(\mathcal{Z}_k^{(r)}) \right] \right\} \quad (17)$$

At the final iteration L (with $N_L = \sum_{l=1}^L n_l$), we reuse all past samples to get the final estimator of P :

$$\widehat{P}_{N_L} = \frac{1}{N_L} \sum_{q=1}^L \sum_{k=1}^{n_l} \mathbb{1}_{\mathcal{Z}_k^{(l)} \in \mathcal{D}} \frac{\pi_0(\mathcal{Z}_k^{(l)})}{q_{\alpha^{(l)}}(\mathcal{Z}_k^{(l)})} \quad (18)$$

Asymptotic optimality and confidence interval. Using theorems 2 and 3 from [6], we can determine sufficient criteria to ensure the consistency and asymptotic normality of the estimator eq. (18).

Theorem 1. *If \mathbb{A} is a compact set, if $\alpha_{opt} \in \mathbb{A}$ is the only minimizer of $\mathcal{D}_{KL}(q_{opt} \| q_{\alpha})$ and if:*

1. *the functions λ , K , and $(U_{\alpha})_{\alpha \in \mathbb{A}}$ are bounded from below and above on their support by positive constants,*
2. *there is $t_{\varepsilon} > 0$ such that $t_z^{\partial} \geq t_{\varepsilon}$ for any $z^- \in \partial E$ and any $z \in \text{supp } K(z^-, \cdot)$,*

then, with $V(\alpha) = \mathbb{E}_{\pi_0} \left[\mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \frac{\pi_0(\mathcal{Z})}{\pi_{\alpha}(\mathcal{Z})} \right] - P^2$, we have :

$$\alpha^{(L)} \xrightarrow[L \rightarrow \infty]{a.s.} \alpha_{opt} \quad \text{and} \quad \sqrt{N_L} \left(\widehat{P}_{N_L} - P \right) \xrightarrow[L \rightarrow \infty]{d} \mathcal{N}(0, V(\alpha_{opt})). \quad (19)$$

The asymptotics can therefore be taken either in the number of iterations L or in the size of the last two samples n_{L-1} and n_L . These are two different yet specific ways to make the total number of simulated trajectories tend towards infinity. The proof of this theorem can be found in [4]. We can also propose a consistent estimator of the asymptotic variance $V(\alpha_{opt})$:

$$\widehat{\sigma}_{N_L}^2 = \frac{1}{N_L} \sum_{q=1}^L \sum_{k=1}^{n_l} \mathbb{1}_{\mathcal{Z}_k^{(l)} \in \mathcal{D}} \frac{\pi_0(\mathcal{Z}_k^{(l)})^2}{q_{\alpha^{(l)}}(\mathcal{Z}_k^{(l)})^2} - \widehat{P}_{N_L}^2. \quad (20)$$

It follows that an asymptotic confidence interval for P with the conditions of theorem 1 is given by:

$$\mathbb{P} \left(P \in \left[\widehat{P}_{n_l} - q_{1-a/2} \widehat{\sigma}_{n_l} n_l^{-1/2}; \widehat{P}_{n_l} + q_{1-a/2} \widehat{\sigma}_{n_l} n_l^{-1/2} \right] \right) \longrightarrow 1 - a, \quad (21)$$

where $q_{1-a/2}$ is the $(1 - a/2)$ -quantile of the $\mathcal{N}(0, 1)$ distribution.

6 NUMERICAL EXPERIMENT

Our adaptive importance sampling (AIS) method is applied to the spent fuel pool system with the parameter values given in section 2 and compared to the CMC method.

Implementation choices.

- The optimization routine called at each step of the cross-entropy procedure to minimize eq. (17) is the method `minimize(·, method=BFGS)` from the `scipy.optimize` toolbox in Python (the gradient of the objective function is explicitly known).
- The AIS method is initialized with a vector α whose starting values are all identical and chosen so that the probability that at least one component failure occurs before the end of the simulation is greater than $1/3$.
- At each iteration, we generate trajectories until we have $n_{\text{CE}} = 10$ faulty trajectories before updating α for $N \in \{10^2, 10^3\}$ and $n_{\text{CE}} = 50$ for $N = 10^4$. We stop the algorithm when the budget N is reached.

Results. We compare the CMC method and the AIS method for different sample sizes in table 3. Looking at the confidence intervals, we see that the AIS method requires between 1000 and 10,000 times less simulations than CMC method to estimate a probability of the order of 10^{-5} with the same level of accuracy.

| Method | N | Estimated probability \hat{P} | Coeff. of var. | 95% confidence interval |
|--------|--------|---------------------------------|----------------|--|
| CMC | 10^5 | 2×10^{-5} | 223.60 | $[0; 4.77 \times 10^{-5}]$ |
| | 10^6 | 1.3×10^{-5} | 277.35 | $[5.93 \times 10^{-6}; 2.01 \times 10^{-5}]$ |
| | 10^7 | 1.77×10^{-5} | 237.68 | $[1.51 \times 10^{-5}; 2.03 \times 10^{-5}]$ |
| AIS | 10^2 | 2.18×10^{-5} | 4.69 | $[1.76 \times 10^{-5}; 4.18 \times 10^{-5}]$ |
| | 10^3 | 2.19×10^{-5} | 3.01 | $[1.78 \times 10^{-5}; 2.60 \times 10^{-5}]$ |
| | 10^4 | 1.99×10^{-5} | 1.01 | $[1.96 \times 10^{-5}; 2.03 \times 10^{-5}]$ |

Table 3: **Comparison of CMC and AIS on the SFP case.** The coefficient of variation is $\hat{\sigma}_N / \hat{P}_N$ with $\hat{\sigma}_N$ from eq. (20) and \hat{P}_N from eq. (18).

Since the confidence interval is constructed from the plug-in estimator of the variance, it is preferable to ensure the stability of our method. In fig. 3, we present 50 different realizations of a confidence interval obtained with the AIS method for a sample size of $N = 10^3$, and compare them to the confidence interval obtained with the CMC method for a sample size of $N = 10^7$. Once again, we observe that most of the intervals obtained by the AIS method outperform the CMC interval. Even the worst interval obtained by AIS returns a range of $[1, 3.5] \times 10^{-5}$, which is already a good estimate for a sample of this size. The AIS method therefore appears to be highly stable.

7 CONCLUSION

Our methodology provides a comprehensive approach for assessing the reliability of hybrid dynamic industrial systems. Since the probability of system failure is typically very small, its crude Monte Carlo estimation requires the simulation of a large number of trajectories of the piecewise deterministic Markov process modeling the system. We have proposed an adaptive importance sampling method to estimate this failure probability accurately while simulating as few PDMP trajectories as possible.

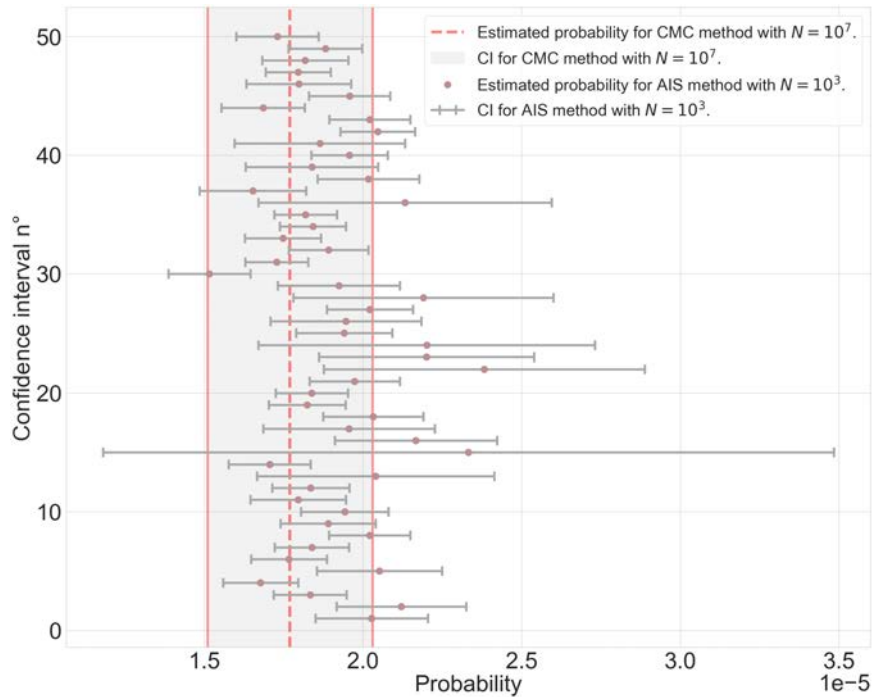


Figure 3: **Comparison of 95% confidence intervals.** Under conditions of table 3, 50 confidence intervals were drawn from the AIS method with $N = 10^3$, and 1 confidence interval was drawn from the CMC method with $N = 10^7$.

We have emphasized in this work the relevance of the committor function to produce efficient importance sampling estimators. We have proposed a parametric family of approximations for the committor function based on the decomposition of the system into minimal path sets (MPS). Our AIS algorithm is based on a cross-entropy procedure with a recycling scheme of past samples to sequentially improve the importance distribution. We have given guarantees for the convergence and the asymptotic normality of the estimator.

We have compared the efficiency of our AIS method to a standard CMC method on a test case from nuclear industry: the spent fuel pool. The AIS method has outperformed the CMC method, reducing the variance by a factor of about 10^4 . We have run the AIS method 50 times, and all the confidence intervals produced were good, indicating that the method is robust and reliable. Moreover, the AIS method has a short "warm-up time" and can produce accurate estimators with sample sizes as small as 10^2 . These features give the AIS method an advantage over the CMC method, even in situations where the occurrence probability P is moderately small (between 10^{-2} and 10^{-4}).

Overall, our methodology provides a powerful tool for estimating the reliability of hybrid dynamic industrial systems, and we recommend its use in all scenarios.

REFERENCES

- [1] Dufour, F., de Saporta, B., & Zhang, H. (2015). *Numerical methods for simulation and optimization of piecewise deterministic Markov processes: application to reliability*. John Wiley & Sons.

- [2] De Boer, P. T., Kroese, D. P., Mannor, S., & Rubinstein, R. Y. (2005). A tutorial on the cross-entropy method. *Annals of operations research*, 134(1), 19-67.
- [3] Chraïbi, H., Dutfoy, A., Galtier, T., & Garnier, J. (2019). On the optimal importance process for piecewise deterministic markov process. *ESAIM: Probability and Statistics*, 23, 893-921.
- [4] Chennetier, G., Chraïbi, H., Dutfoy, A., & Garnier, J. (2022). Adaptive importance sampling based on fault tree analysis for piecewise deterministic Markov process. *arXiv preprint arXiv:2210.16185*.
- [5] Ruijters, E., & Stoelinga, M. (2015). Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools. *Computer science review*, 15, 29-62.
- [6] Portier, F., & Delyon, B. (2018). Asymptotic optimality of adaptive importance sampling. *Advances in Neural Information Processing Systems*, 31:3134-3144.
- [7] Chraïbi, H., Houbedine, J. C., & Sibler, A. (2016). Pycatshoo: Toward a new platform dedicated to dynamic reliability assessments of hybrid systems. In *Proceedings of 13th international conference on probabilistic safety assessment and management (PSAM)*, Seoul, South Korea, October 2016.